Fluid Modeling of Fusion Plasmas with M3D-C1

N.M. Ferraro,[1,2] S.C. Jardin,[3] M.S. Shephard,[4] A. Bauer,[4] J.A. Breslau,[3] J. Chen,[3] F. Delalondre,[4] X. Luo,[4] and F. Zhang[4]

[1]General Atomics, La Jolla, CA
[2]Oak Ridge Institute for Science and Education, Oak Ridge, TN
[3]Princeton Plasma Physics Laboratory, Princeton, NJ
[4]Rensselaer Polytechnic Institute, Troy, NY

E-mail: ferraro@fusion.gat.com

Abstract. The M3D-C1 project is an excellent example of a successful collaboration between three SciDAC centers that has led to a unique computational capability that will be used by the fusion community for years to come. Experience from the CEMM team with related codes let to a particular representation for the magnetic and velocity vector fields that have desirable properties and to a concept for a fully implicit time advance. These considerations led to the requirement and implementation of a 3D finite element with continuous first derivatives. As part of the ITAPS center, the SCOREC team developed customized 3D adaptive gridding and partitioning software appropriate for describing resistive instabilities in toroidal geometry. The TOPS center proposed a concept for an efficient sparse matrix solver in toroidal geometry that consists of a block-Jacobi preconditioner that is motivated by the requirements of strongly magnetized plasma in toroidal geometry. The combined software is shown to be very accurate and efficient, and to have excellent parallel scaling properties. We present initial results of this code applied to internal (sawtooth) and edge (ELMs) instabilities in tokamaks.

1. Introduction

Modeling fusion plasmas is challenging because of the widely disparate scales of length and time characteristic of relevant physical processes. For some applications, these disparities may be eliminated by ordering very fast or slow timescales out of the model equations. This is done in deriving ideal magnetohydrodynamics (MHD), for example, by neglecting slow viscous and resistive processes, very fast electromagnetic waves, and small-scale gyro-motions of particles around the magnetic field. The resulting model is useful for calculating relatively fast motions of the plasma associated with magnetic (Alfvénic) and acoustical waves and has been successful in explaining many gross stability properties of magnetized plasmas. However, ideal MHD is not appropriate for calculating important classes of physical phenomena—in particular, those involving magnetic reconnection, thermal transport, and electrical current equilibration. Furthermore, even fundamentally ideal phenomena often are

strongly influenced by nonideal processes. This is often the case for edge localized modes (ELMs), for example, which will be described in Section 3.

For addressing non-ideal phenomena such as reconnection, or phenomena involving both Alfvénic and transport timescales, a more comprehensive model is required. The M3D-C1 code implements a general two-fluid model, in which the electron and ion species are treated as separate fluids. The model considered in the calculations presented here is as follows, in normalized units:

$$\partial_t n + \nabla \cdot (n\vec{V}) = 0 \tag{1}$$

$$n(\partial_t \vec{V} + \vec{V} \cdot \nabla \vec{V}) = \vec{J} \times \vec{B} - \nabla p - \nabla \cdot \Pi \tag{2}$$

$$\partial_t p + \vec{V} \cdot \nabla p = \Gamma p \nabla \cdot \vec{V} + (\Gamma - 1)(\eta J^2 - \Pi : \nabla \vec{V} - \nabla \cdot \vec{q}) \tag{3}$$

$$- \frac{d_i}{n}\left(\Gamma p_e \frac{\nabla n}{n} - \nabla p_e\right) \tag{4}$$

$$\partial_t B = - \nabla \times \vec{E} \tag{5}$$

where the dynamical fields re n, $\vec{V}$, $\vec{B}$, and p, respectively, the particle number density, fluid velocity, magnetic field, and total pressure. In the present applications, the electron pressure p_e is assumed to be a constant fraction of the total pressure, although M3D-C1 has the capability to advance the electron pressure separately if so desired. The fluid velocity is the full momentum density of the plasma divided by the mass density, which is essentially the velocity of the ion species. The electron velocity is eliminated in favor of the current density, $\vec{J}$, *via* $\vec{V}_e = \vec{V} - d_i\vec{J}/n$. The ratio of specific heats, Γ, is generally given the value 5/3. Here, d_i is the collisionless ion skin depth. In most plasma models, including here, the assumption of quasineutrality is assumed to hold, which means that the ion and electron number densities are equal (n). The electric field $\vec{E}$ and current density $\vec{J}$ are determined by the electron momentum equation and Ampère's law, respectively.

$$\vec{E} = - \vec{V} \times \vec{B} + \eta J + \frac{d_i}{n}(\vec{J} \times \vec{B} - \nabla p) \tag{6}$$

$$\vec{J} = \nabla \times \vec{B} \tag{7}$$

The viscosity comprises isotropic and gyroviscous parts

$$\Pi = \Pi^\circ + \Pi^\wedge \tag{8}$$

$$\Pi^\circ = \mu \mathcal{W} \tag{9}$$

$$\Pi^\wedge = d_i \frac{1}{4}\left[\hat{b} \times \mathcal{W} \cdot (\mathcal{I} + 3\hat{b}\hat{b}) - (\mathcal{I} + 3\hat{b}\hat{b}) \cdot \mathcal{W} \times \hat{b}\right], \tag{10}$$

where $\mathcal{W} = [\nabla \vec{V} + (\nabla \vec{V})^t] - (2/3)\mathcal{I}\nabla \cdot \vec{V}$ is the rate-of-strain tensor, $\hat{b} = \vec{B}/|\vec{B}|$, and μ is an arbitrary scalar field. The heat flux comprises isotropic and parallel diffusivities

$$\vec{q} = - \kappa^\circ \nabla T - \kappa^\parallel \hat{b}\hat{b} \cdot \nabla T, \tag{11}$$

where $T = p/n$ and κ° and $\kappa^\parallel$ are arbitrary scalar fields. All magnetic fusion confinement concepts rely on the fact that thermal conductivity is much larger along magnetic field lines than across them: $\kappa^\parallel \gg \kappa^\circ$.

This model contains a broad range of physical effects and scales. Sophisticated numerical methods are required in order to solve these equations efficiently. To this end, M3D-C1 employs a variety of methods, including implicit and semi-implicit time stepping, high-order elements, and an unstructured mesh. These methods, described in the following section, depend on several software packages being developed in varying degrees together with M3D-C1: the Portable, Extensible Toolkit for Scientific Computation (PETSc) [1, 2, 3], which provides a framework for a number of linear algebraic solvers, including SuperLU [4] (both PETSc and SuperLU are part of the Toward Optimal Petascale Simulation (TOPS) consortium [5]); and adaptive meshing software from the Scientific Computation Research Center (SCOREC), which is part of the Interoperable Tools for Advanced Petascale Simulations (ITAPS) center [6]. By leveraging the powerful methods provided by these packages, M3D-C1 has been able to perform calculations of diverse physical phenomena in realistic physical regimes that previously had not been possible. In Sections 3 and 4, the results of challenging linear and nonlinear calculations of phenomena of significant importance to tokamak experiments are presented.

2. Numerical Methods

2.1. Time Step

Two time-stepping methods are implemented in M3D-C1, each appropriate to different applications. The first is the θ-implicit method, which is the fully implicit method obtained by evaluating each field at the θ-advanced time: $\phi \rightarrow \theta\phi^{n+1} + (1-\theta)\phi^n$. The choice $\theta = 1/2$ obtains the standard Crank-Nicolson method. The other is a split timestep method, in which the velocity and magnetic field are each separately advanced, in that order. The momentum equation, which is solved first, must be solved without knowledge of the magnetic field and pressure at the advanced time. By using equations (5) and (3) to eliminate the occurrences of the advanced-time magnetic field and pressure in the θ-advanced momentum equation, the following discretized form of the momentum equation is obtained [7] (some minor terms are excluded here for clarity):

$$(1-\theta^2\delta t^2\mathcal{L})\vec{V}^{n+1} = (1-\theta^2\delta t^2\mathcal{L})\vec{V}^n + \delta t\left[\vec{J}\times\vec{B} - \nabla p\right]^{n+1/2}, \qquad (12)$$

where $\mathcal{L}$ is the ideal MHD operator [8]. The advanced-time velocity obtained from equation (12) is then used in independent implicit calculations of the density, pressure, and magnetic field (by standard θ-implicit discretization). This method has the advantage that it halves the rank of the matrix equations being solved (although twice as many solves must be performed). For some types of calculations, this gives an accurate result at a fraction of the cost (both in terms of time and memory usage) of the fully-implicit method [7]. For smaller systems in which memory usage and matrix conditioning are less of a concern, the unsplit method is often preferable, as it generally has better convergence properties (although both are technically second-order accurate when θ is chosen appropriately).

2.2. Finite Elements

The M3D-C1 code discretizes poloidal planes using an unstructured These elements possess C^1 continuity: each basis function is both continuous and differentiable everywhere (including across element boundaries). For axisymmetric calculations, a single poloidal plane is considered.

For linear calculations about an axisymmetric equilibrium, the toroidal Fourier components of the linear perturbations are all independent and may be calculated independently. In this case, we represent each field ϕ as

$$\phi(R, \varphi, Z, t) = \tilde{\phi}(R, Z, t)e^{in\varphi},$$

where $\tilde{\phi}(R, Z, t)$ is a complex field. By linearizing the model equations, a system of equations for the complex fields that is independent of φ is obtained. Thus, each complex toroidal Fourier component of the linear solution may be obtained by using a single poloidal plane.

For fully nonlinear 3D calculations, the full torus must be discretized. In M3D-C1, the finite elements are extended in the toroidal direction by taking the tensor product of the reduced quintic elements in a poloidal plane with cubic Hermite elements in the toroidal direction. This yields "brick" elements in the shape of triangular prisms. Because Hermite elements are also C^1, fields remain differentiable in all directions. The use of Hermite elements results in a block-cyclic-tridiagonal matrix structure, in which the diagonal and off-diagonal blocks represent couplings within a plane and between adjacent planes, respectively. For most applications, the coupling within planes will be strong, and therefore the matrix will be block-diagonal dominant. This motivates a block-Jacobi preconditioning strategy, in which the diagonal blocks are factored either completely or incompletely (using SuperLU_dist [4], for example), and then the preconditioned matrix is solved by using an iterative method such as GMRES. In fact, this particular method is now supported by developmental versions of PETSc.

2.3. Flux/Potential Representation

In fluid or gas dynamics, compression of the fluid significantly raises the internal energy, and therefore the most unstable modes of the system are generally incompressible. If a discretization method is not capable of representing exactly incompressible motions of the fluid then unstable solutions may be missed unless the solution is highly resolved. In magnetic fusion plasmas, the compression magnetic field provides an even greater restoring force than compression of the plasma. Writing the velocity as

$$\vec{V} = R^2\nabla U \times \nabla\varphi + R^2\Omega\nabla\varphi + R^{-2}\nabla_\perp\chi, \tag{13}$$

where $\nabla_\perp\chi = (\partial_R\chi\nabla R + \partial_Z\chi\nabla Z)$, is advantageous because neither of the first two terms result in a compression of the toroidal component of the field, and therefore motions that do not compress the toroidal component of the field (which is the dominant component in tokamaks) are exactly representable [9]. The individual terms in equation (13) are

orthogonal under volume integration, which ensures that this representation is never singular.

Representing the magnetic field by using a vector potential

$$\vec{A} = R^2 \nabla\varphi \times \nabla f + \psi \nabla\varphi \tag{14}$$

$$\vec{B} = \nabla \times \vec{A} \tag{15}$$

$$= \nabla\psi \times \nabla\varphi + R^2 \nabla_\perp^2 f \nabla\varphi - \nabla_\perp(\partial_\varphi f) \tag{16}$$

has two advantages. First, the magnetic field manifestly satisfies $\nabla \cdot B = 0$; not only does this enforce an important physical constraint, but it uses that constraint to eliminate a degree of freedom (i.e., the magnetic field is now described, without loss of generality, by only two scalar fields). Second, in the axisymmetric limit, the convenient form

$$\vec{B} = \nabla\psi \times \nabla\varphi + F\nabla\varphi \tag{17}$$

is recovered, where $F = R^2\nabla_\perp^2 f$. This form is useful both because magnetic surfaces are surfaces of constant ψ and because in an ideal-MHD equilibrium $F = F(\psi)$.

Subsets of the scalar fields in the flux/potential representation defined by equations (13) and (16) yield physically meaningful models. For example, advancing only U and ψ yields the "reduced MHD" model, and advancing U, ψ, Ω, and f yields the "four-field" model, each appropriate in certain physical limits [10]. This is a particularly useful property of this representation because it allows the implementation of these reduced models in M3D-C1 with simple modifications to the code.

A potential disadvantage of this representation is that it results in high-order derivatives on the scalar fields. For example, the current density is given by equation (7):

$$\vec{J} = \nabla(R^2\nabla^2 f) \times \nabla\varphi - \Delta^*\psi\nabla\varphi + R^{-2}\nabla_\perp(\partial_\varphi\psi) \tag{18}$$

where $\Delta^*\psi = R^2\nabla_\perp \cdot (\nabla_\perp\psi/R^2)$. Equation (18) involves three derivatives on f. For C^0 elements, which permit only second-order weak derivatives, this equation would have to be broken into two:

$$F^* = R^2\nabla^2 f \tag{19}$$

$$\vec{J} \ = \nabla F^* \times \nabla\varphi - \Delta^*\psi\nabla\varphi + R^{-2}\nabla_\perp(\partial_\varphi\psi), \tag{20}$$

thereby increasing the rank of the matrix. The use of C^1 elements, for which weak derivatives of up to fourth order are permitted, obviates this concern. In this way, C^1 elements allow equations using the flux/potential representation to be implemented in a fundamentally more compact form than C^0 elements.

3. Edge Localized Modes

Edge localized modes (ELMs) present a serious concern for tokamak fusion reactors. These modes are periodic discharges of particles and energy from the plasma edge, thought to be caused by large pressure gradients and current densities present in the edge of high-confinement tokamak plasmas (so-called peeling-ballooning modes [11]). Scaled up to reactor-sized devices, ELMs could quickly damage the plasma-facing components

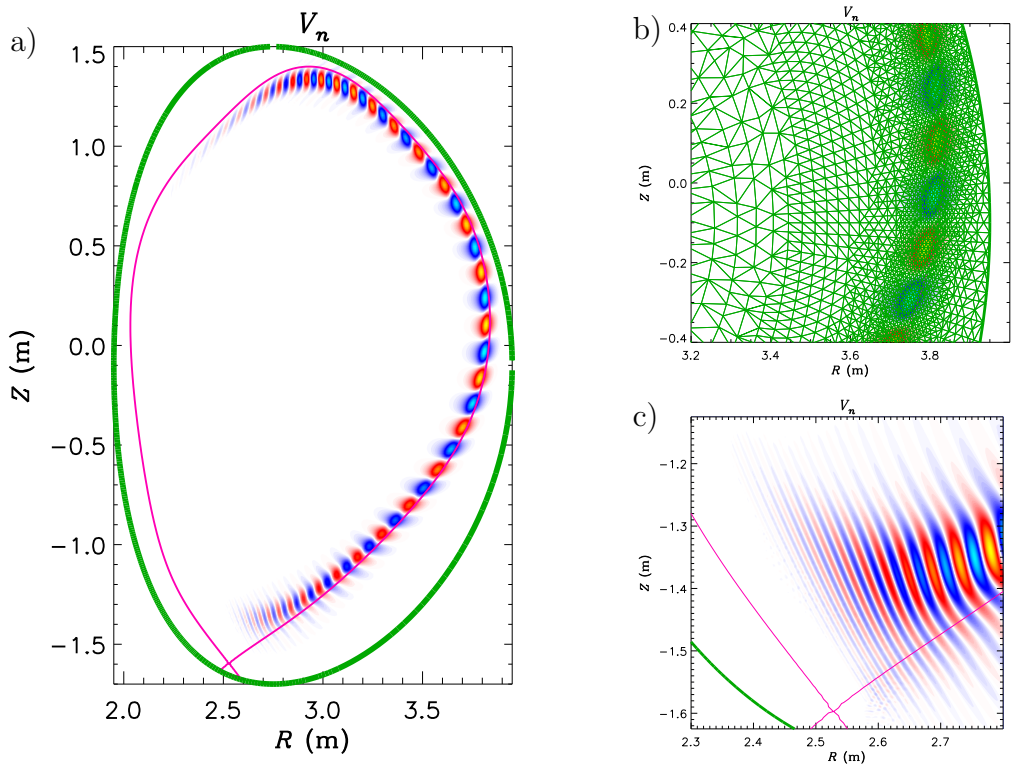

Figure 1. (a) Velocity eigenfunction of the fastest growing peeling-ballooning mode with toroidal mode number $n = 10$. The pink curve represents the plasma-vacuum boundary (the separatrix), and the thick green curve shows the simulation domain boundary. (b) The mesh is packed around the plasma edge in order to resolve the eigenfunction efficiently. (c) The eigenfunction becomes near-singular at the X-point of the separatrix, causing challenges for spatial resolution.

of the device [12]. Understanding ELMs and how to control them is therefore a high-priority issue in magnetic fusion research.

The stability of peeling-ballooning modes is challenging to calculate for two reasons: the mode is localized radially to a narrow region at the edge of a plasma (see Figure 1a); and the eigenmode has very fine structure poloidally. This problem is exacerbated in standard "diverted" tokamak geometries which possess a magnetic separatrix, near which the peeling-ballooning eigenmode becomes almost singular (see Figure 1c). Only in the past decade have specially designed ideal MHD codes been able to calculate peeling-ballooning growth rates accurately [13]. Although ELM stability correlates well with ideal calculations, it is often necessary to include nonideal diamagnetic effects which stabilize the high-n peeling-ballooning modes in order to get good quantitative agreement. More recently, nonideal calculations of peeling-ballooning modes have been performed in simplified geometry [14, 15]. Using M3D-C1, researchers have carried out nonideal calculations of peeling-ballooning modes in full diverted geometry [16]. In these linear calculations, the envelope of the eigenmode is roughly known *a priori*, and therefore a mesh packed to resolve this region greatly increases the efficiency of

the calculation (see Figure 1b). M3D-C1's use of (R, φ, Z) coordinates allows the computational domain to extend across the separatrix, on which a coordinate singularity occurs when a field-aligned coordinate system is used. This is particularly useful when studying ELMs, in which a significant portion of the eigenfunction may cross the separatrix. This also allows M3D-C1 to treat the region outside the separatrix as a resistive plasma rather than a vacuum.

Detailed comparisons between M3D-C1 and ideal MHD codes yielded good agreement in the ideal limit given sufficient resolution. Artificial constraints imposed by ideal models were then relaxed in M3D-C1, providing insight into the roles of resistivity, rotation, and diamagnetic effects on ELM stability. In particular, it was found that treating the cool plasma between the separatrix and the wall as a vacuum, as is usually done in ideal MHD codes, is often a poor approximation, and can lead to a significant change in the growth rate [16]. It was also found that diamagnetic stabilization observed in self-consistent M3D-C1 calculations is generally less effective than simple theory suggests [16]. These calculations may also serve to refute the common misapprehension that resistive codes are unable to perform calculations using resistivities small enough to be relevant to fusion plasmas; resistivities used in these M3D-C1 calculations were equal to or less than those present in modern tokamaks.

4. Sawtooth Cycle

While linear calculations can address the stability or response of given equilibria, nonlinear calculations are required to study the process of a slowly evolving equilibrium crossing a stability threshold. The sawtooth cycle is one of many such processes present in tokamaks. As the plasma slowly heats, the electrical current distribution narrows until it becomes unstable to a kink mode. The instability rapidly flattens the temperature and current density in the core of the plasma, and then the process repeats. (This periodic slow rise and rapid collapse of the core temperature is reason for the name "sawtooth.") Depending on the size of the crashes, this process may be benign or detrimental, but a predictive theory of the expected size of the crashes does not exist. The primary difficulty in simulating the sawtooth cycle is the presence and interaction of the disparate timescales of the heating and instability. Although Alfvénic physics is of little importance through much of the heating phase, it is crucial that this physics be included in the model in order to obtain the instability. Implicit time-stepping methods such as those used by M3D-C1 are highly advantageous for this application, so that timesteps are not limited to stability constraints imposed by the presence of Alfvénic physics.

The first large-scale 3D nonlinear calculations undertaken with M3D-C1 have been simulations of an inductively driven plasma that should exhibit sawtooth crashes. Indeed, these simulations show the plasma quickly settling into a limit cycle characteristic of sawtooth cycles. The M3D-C1 simulation was run for over 20,000 Alfvén periods, which corresponds to roughly 9.5 ms, with each time step spanning 20 Alfvén

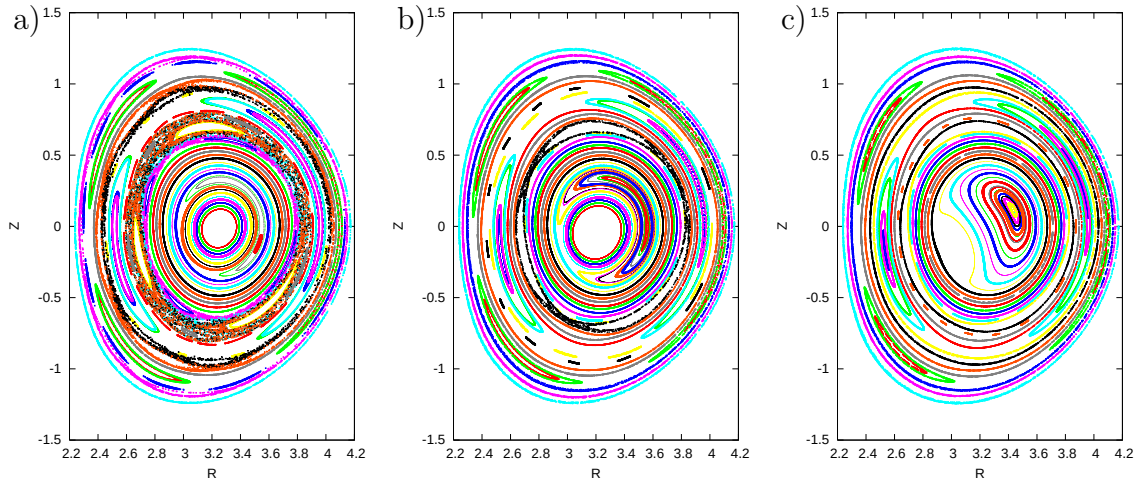

Figure 2. Poincaré plots showing the magnetic field structure (a) before, (b) during, and (c) at the end of a sawtooth crash, as calculated by M3D-C1. Times correspond to $\tau = 0$, 275, and 550 μs, respectively, where $\tau = 0$ corresponds roughly to 2.2 ms from the beginning of the simulation.

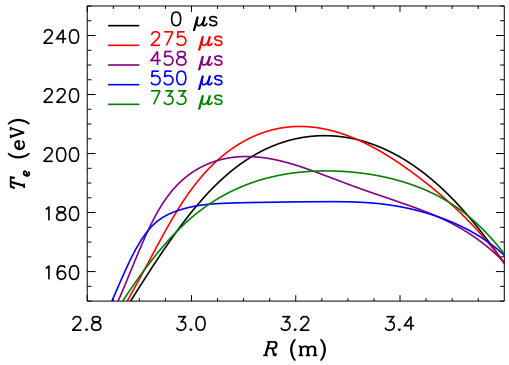

Figure 3. The electron temperature through the midplane ($Z = 0$) for a series of times during a sawtooth cycle, as calculated by M3D-C1. The temperature is peaked at the center ($\tau = 0$) before an instability arises, which causes the temperature profile to skew (τ = 275–458 μs) and flatten (τ = 550 μs). The temperature profile then slowly begins to recover ($\tau = 733$ μs).

periods (for comparison, the Courant limit in this relatively low-resolution case would be roughly 0.1 Alfvén periods). Over the course of the simulation, eight sawtooth cycles are observed. At each sawtooth crash, the magnetic axis is displaced and reconnects with the $q = 1$ surface (this is the surface on which magnetic field lines form periodic orbits, where each orbit circles the torus once toroidally and once poloidally). This process can be seen in Figure 2, where the magnetic axis moves downward and to the left (Figure 2b), and is eventually replaced by the magnetic island that has formed above and to the right of the original axis (Figure 2c). Because of the anisotropic

thermal conduction included in these calculations, the isotherms closely align with the magnetic surfaces. This causes the peaked temperature profile to skew and eventually flatten as the magnetic axis shifts reconnects with the $q = 1$ surface, as can be seen in Figure 3. Once flattened, the plasma is again stable, and the slow Joule heating process begins to build the peaked temperature profile again to repeat the cycle.

5. Summary

M3D-C1 implements a comprehensive fluid model that encompasses numerous disparate spatial and temporal scales. The efficient solution of this model requires advanced numerical techniques such as implicit time-steps, high-order finite elements, and nonuniform meshing. Ongoing developments in both solver and meshing packages made available through the TOPS and ITAPS centers are continuing to provide M3D-C1 with state-of-the-art tools for implementing these numerical methods. Furthermore, as part of the CEMM project, M3D-C1 also benefits from close collaboration with the developers of the NIMROD [17] and M3D [18] codes. Each of these codes implements similar models but uses different numerical methods. The range of numerical methods utilized among the various codes is much greater in breadth than would be feasible within a single framework. By comparing code performance for various applications, the most efficient methods for each application may be determined in this way [19, 7].

The combination of efficient spatial discretization and implicit time-stepping in M3D-C1 has enabled simulations of important tokamak phenomena that have previously only been possible using simplified models. Challenging calculations of peeling-ballooning modes in toroidal, diverted geometry and with realistic values of resistivity have been performed, with the results validating and extending calculations using reduced models. Semi-implicit time-stepping methods have also facilitated nonlinear simulations of sawtooth cycles, which span transport timescales. One of the great advantages of M3D-C1 is its general applicability to many phenomena; although not discussed here, M3D-C1 has also been applied to magnetic reconnection in periodic slab geometry, and to the magnetorotational and magnetothermal instabilities present in astrophysical accretion disks and stellar atmospheres [20].

Now that M3D-C1 is reaching maturity as a fully-featured 3D nonlinear code, it can be applied to many other physical applications. In particular, while individual ELM crashes have been studied both linearly and nonlinearly, a repeating ELM cycle as observed in experiments has not been achieved yet with a realistic plasma model; we believe that M3D-C1 would be well suited for this. Furthermore, many observed features of the sawtooth cycle do not agree with simple theory. More extensive M3D-C1 calculations than those presented here are planned and could help to explain these discrepancies. M3D-C1 also is capable of obtaining nonaxisymmetric steady states containing magnetic islands. Such states are frequently observed in tokamaks, usually when an instability grows and saturates nonlinearly, but are difficult or impossible to describe using reduced models. M3D-C1 provides a promising tool for exploring this

important type of phenomenon.

Acknowledgments

We are grateful for many helpful conversations with B. Burke, P. Snyder, and L. Sugiyama regarding ELM calculations, and M. Chance, A. Glasser, X. Li, and C. Sovinec regarding numerical methods.

References

[1] Satish Balay, Kris Buschelman, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley, Lois Curfman McInnes, Barry F. Smith, and Hong Zhang. PETSc Web page, 2001. http://www.mcs.anl.gov/petsc.

[2] Satish Balay, Kris Buschelman, Victor Eijkhout, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley, Lois Curfman McInnes, Barry F. Smith, and Hong Zhang. PETSc users manual. Technical Report ANL-95/11 - Revision 2.1.5, Argonne National Laboratory, 2004.

[3] Satish Balay, William D. Gropp, Lois Curfman McInnes, and Barry F. Smith. Efficient management of parallelism in object oriented numerical software libraries. In E. Arge, A. M. Bruaset, and H. P. Langtangen, editors, *Modern Software Tools in Scientific Computing*, pages 163–202. Birkhäuser Press, 1997.

[4] Xiaoye S. Li and James W. Demmel. SuperLU_DIST: A scalable distributed-memory sparse direct solver for unsymmetric linear systems. *ACM Trans. Mathematical Software*, 29(2):110–140, June 2003.

[5] TOPS Web page, 2008. http://www.scalablesolvers.org.

[6] ITAPS Web page, December 2010. http://www.itaps.org.

[7] N. M. Ferraro and S. C. Jardin. Calculations of two-fluid magnetohydrodynamic axisymmetric steady-states. *J. Comp. Phys.*, 228(1):7742–7770, November 2009.

[8] I. B. Bernstein, E. A. Frieman, M. D. Kruskal, and R. M. Kulsrud. An energy principle for hydromagnetic stability problems. *Proc. Roy. Soc. London. Series A*, 244(1236):17–40, February 1958.

[9] J. Breslau, N. Ferraro, and S. Jardin. Some properties of the M3D-C^1 form of the three-dimensional magnetohydrodynamics equations. *Phys. Plasmas*, 16:092503, September 2009.

[10] S. C. Jardin, J. Breslau, and N. Ferraro. A high-order implicit finite element method for integrating the two-fluid magnetohydrodynamic equations in two dimensions. *J. Comp. Phys.*, 226(2):2146–2174, October 2007.

[11] J. W. Connor, R. J. Hastie, H. R. Wilson, and R. L. Miller. Magnetohydrodynamic stability of tokamak edge plasmas. *Phys. Plasmas*, 5(7):2687–2700, July 1998.

[12] G. Federici, A. Loarte, and G. Strohmayer. Assessment of erosion of the ITER divertor targets during type I ELMs. *Plasma Phys. Control. Fusion*, 45:1523–1547, 2003.

[13] P. B. Snyder, H. R. Wilson, J. R. Ferron, L. L. Lao, W. Leonard, T. H. Osborne, and A. D. Turnbull. Edge localized modes and the pedestal: a model based on coupled peeling-ballooning modes. *Phys. Plasmas*, 9:2037, May 2002.

[14] B. J. Burke, S. E. Kruger, C. C. Hegna, P. Zhu, P. B. Snyder, C. R. Sovinec, and E. C. Howell. Edge localized linear ideal magnetohydrodynamic instability studies in an extended-magnetohydrodynamic code. Submitted to Phys. Plasmas., 2010.

[15] X. Q. Xu, B. Dudson, P. B. Snyder, M. V. Umansky, and H. Wilson. Nonlinear simulations of peeling-ballooning modes with anomalous electron viscsosity and their role in edge localized mode crashes. *Phys. Rev. Lett.*, 105:175005, 2010.

[16] N. M. Ferraro, S. C. Jardin, and P. B. Snyder. Ideal and resistive edge stability calculations with M3D-C^1. *Phys. Plasmas*, 17:102508, 2010.

[17] A. H. Glasser, C. R. Sovinec, R. A. Nebel, T. A. Gianakon, S. J. Plimpton, M. S. Chu, D. D. Schnack, and the NIMROD Team. The NIMROD code: a new approach to numerical plasma physics. *Plasma Phys. Control. Fusion*, 41:A747–A755, 1999.

[18] W. Park, E. V. Belova, G. Y. Fu, X. Z. Tang, H. R. Strauss, and L. E. Sugiyama. Plasma simulation studies using multilevel physics models. *Phys. Plasmas*, 6(5):1796–1803, May 1999.

[19] J. A. Breslau, C. R. Sovinec, and S. C. Jardin. An improved tokamak sawtooth benchmark for 3D nonlinear MHD. *Comm. Comp. Phys.*, 4(3):647–658, September 2008.

[20] Nathaniel M. Ferraro. *Non-ideal effects on the stability and transport of magnetized plasmas.* PhD thesis, Princeton University, November 2008.